Supporting Information

Reddy et al. 10.1073/pnas.0907330106

SI Methods

Experimental Design. As shown in Fig. 1, within each block, the images of a given category were alternately presented to the left and right of fixation. This design prevented subjects from attending to just one side of fixation during an entire block, thus avoiding potential confounds due to hemisphere-specific activations. Although our attentional manipulation clearly entails a "category-specific" component (only one or two object categories were task-relevant at any time), it may also involve spatial attention, because the 800-ms stimulus duration and the predictability of the alternating sequence would have allowed subjects to spatially attend to the relevant location on each stimulus presentation. The images were centered at 4° on either side of fixation and subtended ≈7 × 7° of visual angle. Each subject was tested on six to seven experimental runs in the scanner. The order of blocks was counterbalanced across subjects. Eye tracking was performed in a separate session with an IR IScan Camera, sampling the eye position at 240 Hz. Viewing conditions in the scanner were replicated as closely as possible during these sessions. Subjects were not explicitly instructed to fixate during these sessions but were told to perform the experiment with the same strategy they had used during the fMRI scanning session (when they had been strictly instructed to fixate). As reported in detail in ref. 1, subjects reliably maintained fixation and did not make eye movements to the peripherally presented stimuli.

Projection Analysis. As also described below (see also Fig. 2), our analysis relied on simple mathematical projections in a multidimensional space: Each fMRI activation pattern was associated with a vector in multidimensional space, which was then projected either onto a plane (Fig. 3; Figs. S1 and S2) or onto a line (Fig. 4; Figs. S3 and S4). The plane and line were defined (as described in more detail below) based on two reference vectors, corresponding to isolated presentations of categories X and non-X. For both the plane and line analyses, a leave-one-out procedure was used: The isolated conditions in N − 1 runs were used to define the plane or the weighted average line, respectively, and data from the Nth run was then projected onto this plane or weighted average line. Additionally, for each category X, all analyses were performed pair-wise (e.g., face-house, face-shoe, and face-car) and then averaged to give a single value for (X, non-X).

Plane Projections. The first step of the analysis asked whether the joint response during paired presentation, Pair (X, non-X), corresponds better to a weighted average of the response to the individual stimuli (X, non-X) or to a weighted sum.

That is, if Pair (X, non-X) = α.X + β.non-X + Δ, is $\alpha+\beta \sim 1$ (the weighted average model), or is $\alpha + \beta \sim 2$ (the weighted sum model)?

The responses to X and non-X in the multidimensional voxel space correspond to high-dimensional vectors, and the plane containing these two vectors also contains the weighted average and weighted sum lines (see Fig. 2). Thus, to determine which model better describes the response in the paired presentation condition, the corresponding vector Pair (X, non-X) can be projected onto this plane. The distance of this projected point from the weighted average and weighted sum lines indicates which model better suits the data; the shorter the distance, the better the model.

The projection of a point P to point P_1 on the plane containing vectors X and non-X can be computed by solving the following three equations:

$$\vec{P}_1 = \alpha\cdot\vec{X} + \beta\cdot\overrightarrow{\text{non-X}} \quad [1]$$

$$(\vec{P}_1 - \vec{P})\cdot\vec{X} = 0 \quad [2]$$

$$(\vec{P}_1 - \vec{P})\cdot\overrightarrow{\text{non-X}} = 0 \quad [3]$$

Solving for α and β yields:

$$\alpha = \frac{(\vec{P}\cdot\overrightarrow{\text{non-X}})\cdot(\vec{X}\cdot\overrightarrow{\text{non-X}}) - (\vec{P}\cdot\vec{X})\cdot|\overrightarrow{\text{non-X}}|^2}{(\vec{X}\cdot\overrightarrow{\text{non-X}})\cdot(\vec{X}\cdot\overrightarrow{\text{non-X}}) - |\vec{X}|^2\cdot|\overrightarrow{\text{non-X}}|^2}$$

$$\beta = (\vec{P}\cdot\vec{X} - \alpha\vec{X}\cdot\vec{X})/(\vec{X}\cdot\overrightarrow{\text{non-X}})$$

Note that $|\vec{P}_1 - \vec{P}| = |\vec{\Delta}|$ would correspond to the distance of P from the plane, and would thus be a measure of how much of the response to Pair (X, non-X) is not explained by a simple linear model (see Fig. S2).

Weighted Average Line Projections. The distance of the plane projection P_1 of Pair (X, non-X) from the weighted average and weighted sum lines indicated that the projection error was lower for the weighted average model (Fig. 3; Fig. S1). In the second step of the analysis, we therefore asked, for all paired responses, what values of α and β (i.e., the weights) satisfied the equation: Pair (X, non-X) = α.X + β.non-X + Δ', such that $\alpha+\beta = 1$ for all of the paired responses. Note that the constraint $\alpha+\beta = 1$ implies that the projection $P_2 = \alpha$.X + β.non-X lies along the line joining points X and non-X (i.e., the weighted average line).

To solve this equation for α (and $\beta = 1 - \alpha$), we thus projected Pair (X, non-X) onto the line between points X and non-X, using the following equation:

$$\alpha = (\vec{X} - \overrightarrow{\text{non-X}})\cdot(\vec{P} - \overrightarrow{\text{non-X}})/|\vec{X} - \overrightarrow{\text{non-X}}|^2$$

The value $|\vec{P}_2 - \vec{P}| = |\vec{\Delta}'|$ in this case reflects the variability that the weighted average model cannot account for (Fig. S3). Note that some authors have also argued in favor of a "max" model (2, 3), the projection error for the max model was also computed but turned out to be larger than for the weighted average model, and was thus not further considered in the analysis.

The attentional bias was computed as the ratio of the distance between the "attend-X" (or "attend-nonX") point in Fig. 4 to the "isolated X" (or "isolated non-X") point, and the distance between the "attend both" point to the "isolated-X" (or "isolated-nonX") point. To avoid dividing by zero or by very small numbers (that would artificially inflate the index in either the positive or negative directions), we did not consider instances where the "attend-both" condition was implausibly close to either isolated condition (i.e., if the denominator in the ratio was <0.1).

1. Reddy L, Kanwisher N (2007) Category selectivity in the ventral visual pathway confers robustness to clutter and diverted attention. *Curr Biol* 17:2067–2072.
2. Gawne TJ, Martin JM (2002) Responses of primate visual cortical V4 neurons to simultaneously presented stimuli. *J Neurophysiol* 88:1128–1135.
3. Riesenhuber M, Poggio T (1999) Hierarchical models of object recognition in cortex. *Nat Neurosci* 2:1019–1025.

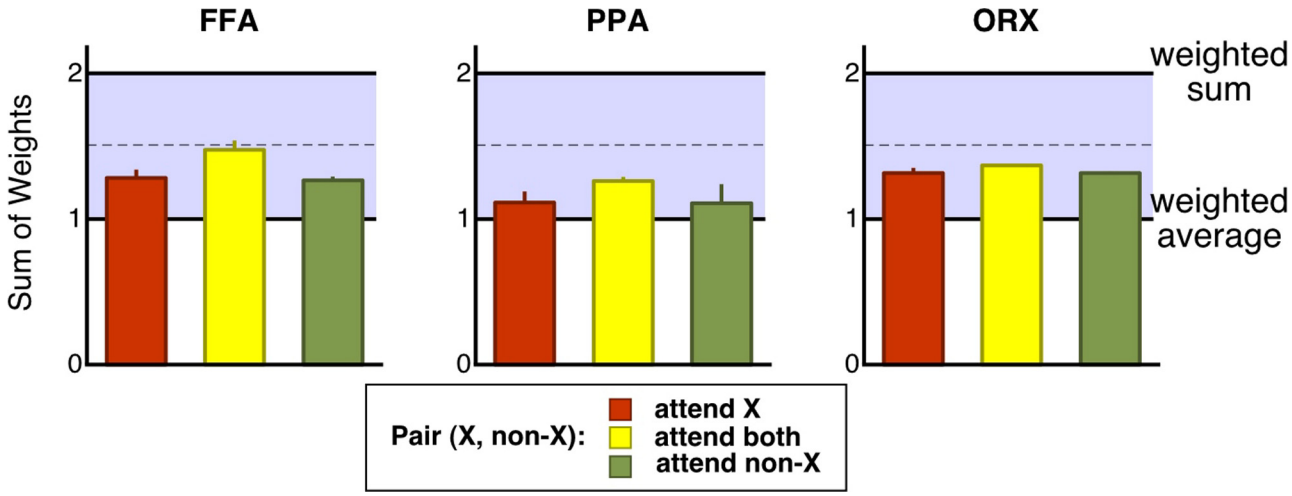

Fig. S1. An index evaluating the distance of each Pair (X, non-X) in Fig. 3 to the expected weighted average and weighted sum lines for the fusiform face area (FFA), parahippocampal place area (PPA), and object responsive voxels in ventral temporal cortex (ORX); see *Methods*. A value of 1 would reflect a true weighted average, and a value of 2 a true weighted sum. The shaded blue area corresponds to the blue area in Fig. 3.

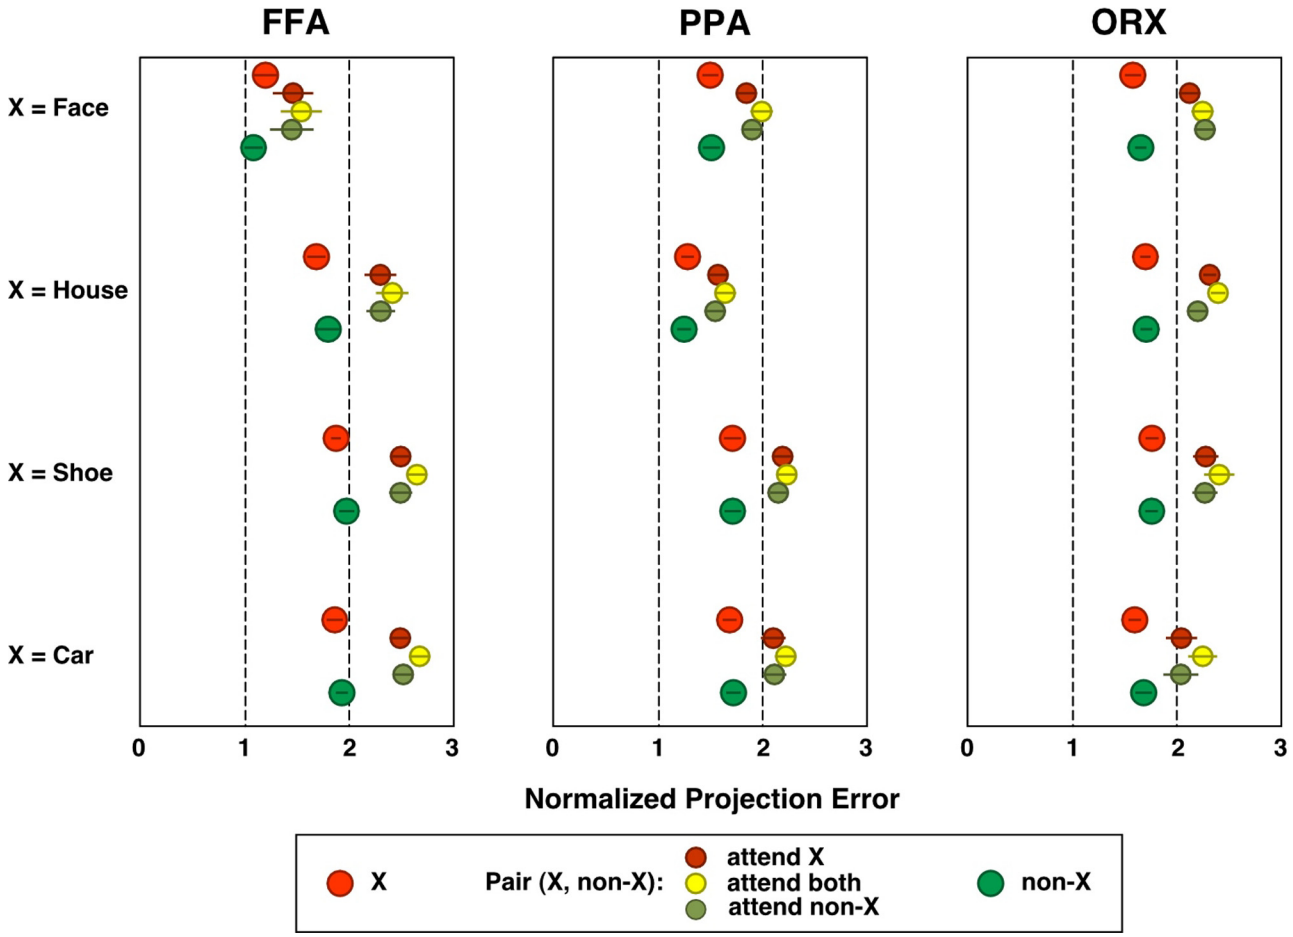

Fig. S2. The normalized projection error quantifying the distance of the response in each condition from the plane defined by the X and non-X vectors (normalized to the norm of vector [X, non-X]). As in Figs. 3 and 4, for each X, this analysis was computed pair-wise over all combinations of (X, non-X) and the averaged results are shown here. SEM values were calculated across subjects. The normalized projection errors for the isolated conditions reflect the block-to-block variability in the dataset (i.e., the responses in different repetitions of a given condition are not identical). The additional normalized projection error values for the pairs (X, non-X), beyond those observed for isolated presentations reflect the variability (and noise) that neither the weighted sum nor weighted average models can account for. The projection error for the paired conditions (X, non-X) is only 34 $\pm$ 2% larger than the corresponding projection error for the isolated conditions. Therefore, the inherent variability and measurement noise that similarly affect the isolated and paired conditions account for the larger portion of the projection error.

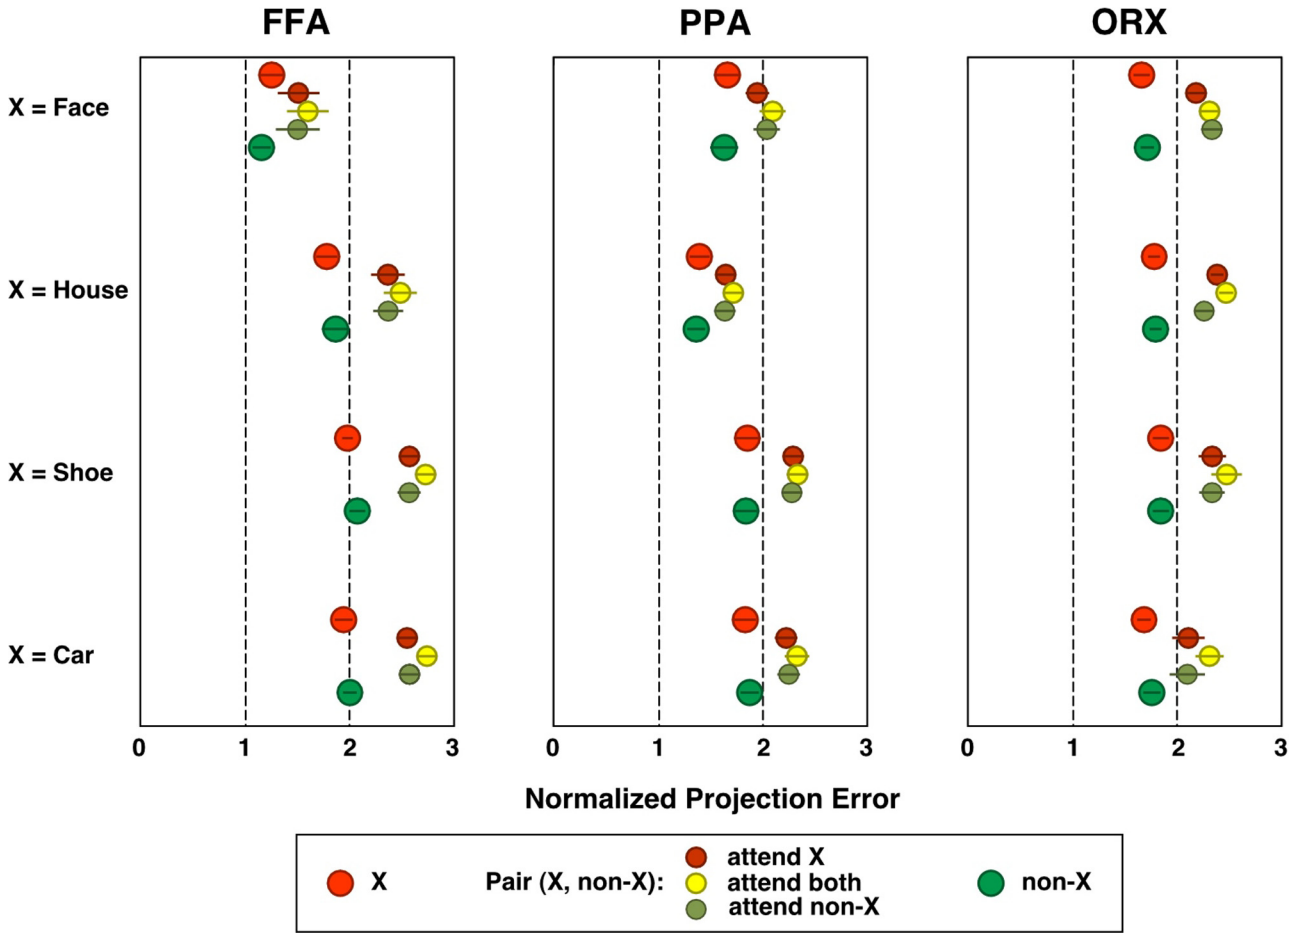

Fig. S3. The normalized projection error quantifying the distance of the response in each condition from the weighted average line defined by the X and non-X vectors (normalized to the distance between points X and non-X). As previously, for each X, this analysis was computed pair-wise over all combinations of (X, non-X) and the averaged results are shown here. SEM values were calculated across subjects. The projection error for the paired conditions (X, non-X) is only 31 ± 2% larger than the corresponding projection error for the isolated conditions. Note that even though Figs. S2 and S3 look quite similar to each other slight differences can be observed. The average normalized projection error is 1.95 in Fig. S2 and 2.02 in Fig. S3. Furthermore, the projection error for each point in Fig. S3 is larger than the error for the corresponding point in Fig. S2 as expected (because the line belongs to the plane, plane projection errors cannot be larger than line projection errors).

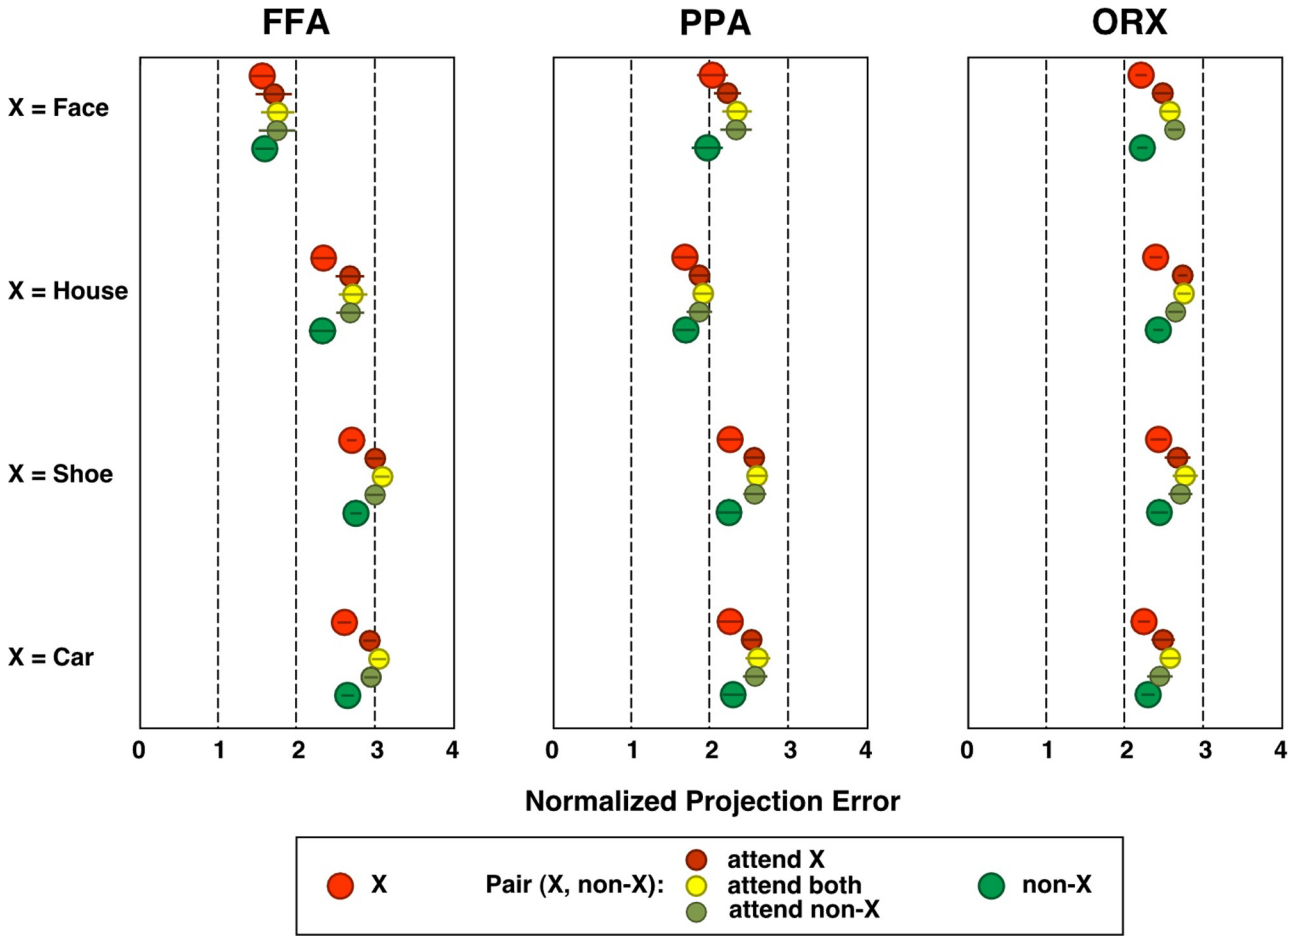

Fig. S4. The normalized projection error quantifying the distance of the response in each condition from the weighted sum line defined by the X and non-X vectors (normalized to the distance between points X and non-X). As previously, for each X, this analysis was computed pair-wise over all combinations of (X, non-X) and the averaged results are shown here. SEM values were calculated across subjects. The averaged normalized projection error from the weighted sum line was larger than the corresponding distance from the weighted average line (compare with Fig. S3; note the difference in the *x*-axis scale).

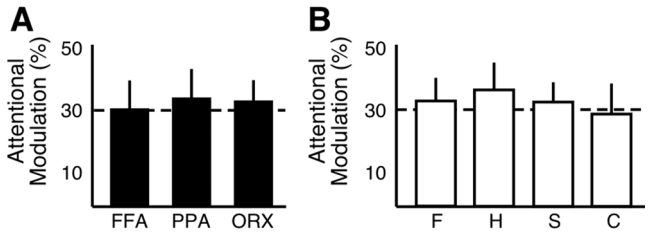

Fig. S5. In a second measure of the attentional bias, we considered the bias to be the distance between the attended conditions normalized by the distance between the isolated conditions on the weighted average line projection. This measure has the advantage of depending only on the projections of the attend-X and attend-nonX conditions (and of the corresponding isolated conditions), but not on the attend-both condition. The results of this analysis are consistent with the attentional bias reported in Fig. 5.

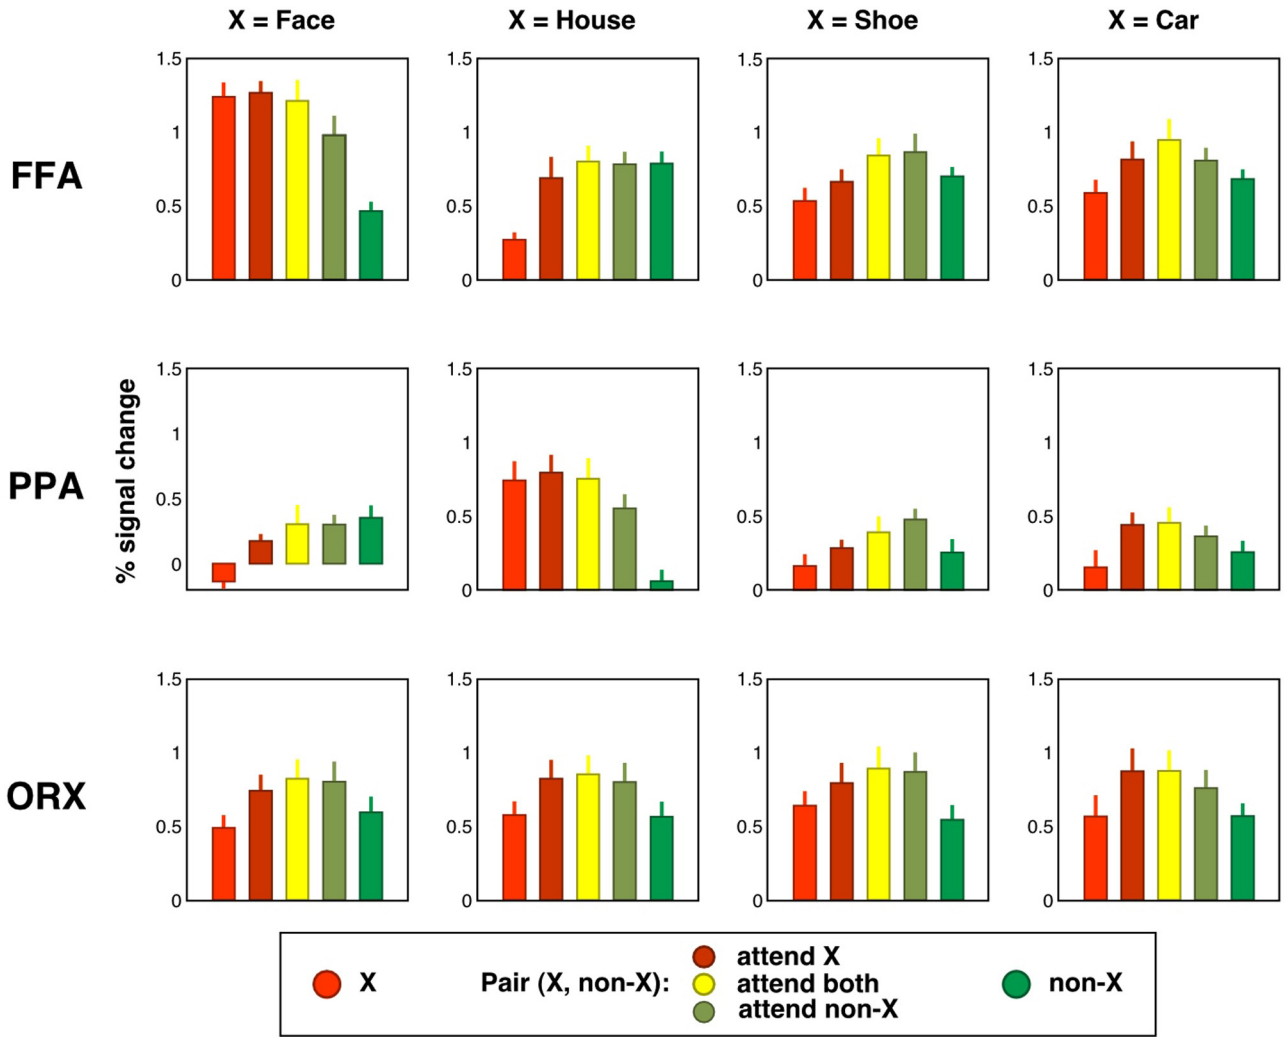

Fig. S6. The blood-oxygen-level-dependent (BOLD) signal change in each region of interest (ROI) for the four categories in the five experimental conditions. Note that the absolute level of BOLD response for isolated stimulus presentations sometimes does not allow one to distinguish between categories (e.g., for all categories in ORX), whereas the multivoxel response does (see Fig. 4). Furthermore, the absolute BOLD response does not reveal a consistent pattern for how multiple stimuli are combined. For instance, in those cases where a ''preferred'' and a ''nonpreferred'' category exist (e.g., for faces and houses in the FFA and PPA), the response to paired presentations tends to fall between the isolated responses (as predicted by a weighted average model). However, when no preferred stimulus is apparent (e.g., shoes and cars in all three ROIs), there appears to be an additive component in the paired response (suggestive of a weighted sum model). In contrast, looking at response combination in the multivoxel space did reveal a consistent pattern, across categories and ROIs, in favor of a weighted average model (Fig. 3; Fig. S1).